

Potentials of mean force for protein structure prediction vindicated, formalized and generalized

Thomas Hamelryck^{1,†,#}, Mikael Borg^{1,#}, Martin Paluszewski^{1,#}, Jonas Paulsen¹, Jes Frellsen¹, Christian Andreetta¹, Wouter Boomsma^{2,3}, Sandro Bottaro², Jesper Ferkinghoff-Borg^{2,‡}

1 Bioinformatics Center, Department of Biology, University of Copenhagen, Copenhagen, Denmark

2 Biomedical engineering, DTU Elektro, Technical University of Denmark, Lyngby, Denmark

3 Department of Chemistry, University of Cambridge, Cambridge, United Kingdom

† E-mail: thamelry@binf.ku.dk

‡ E-mail: jfb@elektro.dtu.dk

Joint first authors.

Abstract

Understanding protein structure is of crucial importance in science, medicine and biotechnology. For about two decades, knowledge based potentials based on pairwise distances – so-called “potentials of mean force” (PMFs) – have been center stage in the prediction and design of protein structure and the simulation of protein folding. However, the validity, scope and limitations of these potentials are still vigorously debated and disputed, and the optimal choice of the reference state – a necessary component of these potentials – is an unsolved problem. PMFs are loosely justified by analogy to the reversible work theorem in statistical physics, or by a statistical argument based on a likelihood function. Both justifications are insightful but leave many questions unanswered. Here, we show for the first time that PMFs can be seen as approximations to quantities that do have a rigorous probabilistic justification: they naturally arise when probability distributions over different features of proteins need to be combined. We call these quantities “reference ratio distributions” deriving from the application of the “reference ratio method”. This new view is not only of theoretical relevance, but leads to many insights that are of direct practical use: the reference state is uniquely defined and does not require external physical insights; the approach can be generalized beyond pairwise distances to arbitrary features of protein structure; and it becomes clear for which purposes the use of these quantities is justified. We illustrate these insights

with two applications, involving the radius of gyration and hydrogen bonding. In the latter case, we also show how the reference ratio method can be iteratively applied to sculpt an energy funnel. Our results considerably increase the understanding and scope of energy functions derived from known biomolecular structures.

Introduction

Methods for protein structure prediction, simulation and design rely on an energy function that represents the protein’s free energy landscape; a protein’s native state typically corresponds to the state with minimum free energy [1]. So-called knowledge based potentials (KBP) are parametrized functions for free energy calculations that are commonly used for modeling protein structures [2, 3]. These potentials are obtained from databases of known protein structures and lie at the heart of some of the best protein structure prediction methods. The use of KBPs originates from the work of Tanaka and Scheraga [4] who were the first to extract effective interactions from the frequency of contacts in X-ray structures of native proteins. Miyazawa and Jernigan formalized the theory for contact interactions by means of the quasi-chemical approximation [5, 6].

Many different approaches for developing KBPs exist, but the most successful methods to date build upon a seminal paper by Sippl – published two decades ago – which introduced KBPs based on probability distributions of pairwise distances in proteins and reference states [7]. These KBPs were called “potentials of mean force”, and seen as approximations of free energy functions. Sippl’s work was inspired by the statistical physics of liquids, where a “potential of mean force” has a very precise and undisputed definition and meaning [8, 9]. However, the validity of the application to biological macromolecules is vigorously disputed in the literature [2, 10–17]. Nonetheless, PMFs are widely used with considerable success; not only for protein structure prediction [3, 18, 19], but also for quality assessment and identification of errors [20–22], fold recognition and threading [23, 24], molecular dynamics [24], protein-ligand interactions [16, 25], protein design and engineering [26, 27], and the prediction of binding affinity [17, 28]. In this article, the abbreviation “PMF” will refer to the pairwise distance dependent KBPs following Sippl [7], and the generalization that we introduce in this article; we will write “potentials of mean force” in full when we refer to the real, physically valid potentials as used in liquid systems [9, 13, 29]. At the end of the article, we will propose a new name for these statistical quantities, to set them apart from true potentials of

mean force with a firm physical basis.

Despite the progress in methodology and theory, and the dramatic increase in the number of experimentally determined protein structures, the accuracy of the energy functions still remains the main obstacle to accurate protein structure prediction [22, 30, 31]. Recently, several groups demonstrated that it is the quality of the coarse grained energy functions [18], rather than inadequate sampling, that impairs the successful prediction of the native state [30, 31]. The insights presented in this article point towards a new, theoretically well-founded way to construct and refine energy functions, and thus address a timely problem.

We start with an informal outline of the general ideas presented in this article, and then analyze two notable attempts in the literature to justify PMFs. We point out their shortcomings, and subsequently present a rigorous probabilistic explanation of the strengths and shortcomings of traditional pairwise distance PMFs. This explanation sheds a surprising new light on the nature of the reference state, and allows the generalization of PMFs beyond pairwise distances in a statistically valid way. Finally, we demonstrate our method in two applications involving protein compactness and hydrogen bonding. In the latter case, we also show that PMFs can be iteratively optimized, thereby effectively sculpting an energy funnel [24, 32–36].

Results and discussion

Overview

In order to emphasize the practical implications of the theoretical insights that we present here, we start with a very concrete example that illustrates the essential concepts (see Fig. 1). Currently, protein structure prediction methods often make use of fragment libraries: collections of short fragments derived from known protein structures in the Protein Data Bank (PDB). By assembling a suitable set of fragments, one obtains conformations that are protein-like on a local length scale. That is, these conformations typically lack non-local features that characterize real proteins, such as a well-packed hydrophobic core or an extensive hydrogen bond network. Such aspects of protein structure are not, or only partly, captured by fragment libraries.

Formally, a fragment library specifies a probability distribution $Q(X)$, where X is for example a vector of dihedral angles. In order to obtain conformations that also possess the desired non-local features, $Q(X)$

needs to be complemented with another probability distribution $P(Y)$, with Y being for example a vector of pairwise distances, the radius of gyration, the hydrogen bonding network, or any combination of non-local features. Typically, Y is a deterministic function of X ; we use the notation $Y(X)$ when necessary.

For the sake of argument, we will focus on the radius of gyration (r_g) at this point; in this case $Y(X)$ becomes $r_g(X)$. We assume that a suitable $P(r_g)$ was derived from the set of known protein structures; without loss of generality, we leave out the dependency on the amino acid sequence for simplicity. The problem that we address in this article can be illustrated with the following question: how can we combine $P(r_g)$ and $Q(X)$ in a rigorous, meaningful way? In other words, we want to use the fragment library to sample conformations whose radii of gyration r_g are distributed according to $P(r_g)$. These conformations should display a realistic *local* structure as well, reflecting the use of the fragment library. Simply multiplying $P(r_g(X))$ and $Q(X)$ does not lead to the desired result, as X and R_g are not independent; the resulting conformations will not be distributed according to $P(r_g)$.

The solution is given in Fig. 1; it involves the probability distribution $Q_R(r_g)$, the probability distribution over the radius of gyration for conformations sampled solely from the fragment library. The subscript R stands for *reference state* as will be explained below. The solution generates conformations whose radii of gyration are distributed according to $P(r_g)$. The influence of $Q(X)$ is apparent in the fact that for conformations with a given r_g , their local structure X will be distributed according to $Q(X | r_g)$. The latter distribution has a clear interpretation: it corresponds to sampling an infinite amount of conformations from a fragment library, and retaining only those with the desired r_g . Note that even if we chose the uniform distribution for $Q(X)$, the resulting $Q_R(r_g)$ will *not* (necessarily) be uniform.

Intuitively, $P(r_g)$ provides correct information about the radius of gyration, but no information about local structure; $Q(X)$ provides approximately correct information about the structure of proteins on a local length scale, but is incorrect on a global scale (leading to an incorrect probability distribution for the radius of gyration); finally, the formula shown in Fig. 1 merges these two complementary sources of information together. Another viewpoint is that $P(r_g)$ and $Q(r_g)$ are used to correct the shortcomings of $Q(X)$. This construction is statistically rigorous, provided that $P(r_g)$ and $Q(X)$ are proper probability distributions.

After this illustrative example, we now review the use of PMFs in protein structure prediction, and discuss how PMFs can be understood and generalized in the theoretical framework that we briefly outlined here.

Pairwise PMFs for protein structure prediction

Many textbooks present PMFs as a simple consequence of the Boltzmann distribution, as applied to pairwise distances between amino acids. This distribution, applied to a specific pair of amino acids, is given by:

$$P(r) = \frac{1}{Z} e^{-\frac{F(r)}{kT}}$$

where r is the distance, k is Boltzmann's constant, T is the temperature and Z is the partition function, with $Z = \int e^{-\frac{F(r)}{kT}} dr$. The quantity $F(r)$ is the free energy assigned to the pairwise system. Simple rearrangement results in the *inverse Boltzmann formula*, which expresses the free energy $F(r)$ as a function of $P(r)$:

$$F(r) = -kT \ln P(r) - kT \ln Z$$

To construct a PMF, one then introduces a so-called *reference state* with a corresponding distribution Q_R and partition function Z_R , and calculates the following free energy difference:

$$\Delta F(r) = -kT \ln \frac{P(r)}{Q_R(r)} - kT \ln \frac{Z}{Z_R} \quad (1)$$

The reference state typically results from a hypothetical system in which the specific interactions between the amino acids are absent [7]. The second term involving Z and Z_R can be ignored, as it is a constant.

In practice, $P(r)$ is estimated from the database of known protein structures, while $Q_R(r)$ typically results from calculations or simulations. For example, $P(r)$ could be the conditional probability of finding the $C\beta$ atoms of a valine and a serine at a given distance r from each other, giving rise to the free energy difference ΔF . The total free energy difference of a protein, ΔF_{TOT} , is then claimed to be the sum of all the pairwise free energies:

$$\Delta F_{\text{TOT}} = \sum_{i < j} \Delta F(r_{ij} | a_i, a_j) \quad (2)$$

$$= -kT \sum_{i < j} \ln \frac{P(r_{ij} | a_i, a_j)}{Q_R(r_{ij} | a_i, a_j)} \quad (3)$$

where the sum runs over all amino acid pairs a_i, a_j (with $i < j$) and r_{ij} is their corresponding distance.

It should be noted that in many studies Q_R does not depend on the amino acid sequence [11].

Intuitively, it is clear that a low free energy difference indicates that the set of distances in a structure is more likely in proteins than in the reference state. However, the physical meaning of these PMFs have been widely disputed since their introduction [2, 12–15]. Indeed, why is it at all necessary to subtract a reference state energy? What is the optimal reference state? Can PMFs be generalized and justified beyond pairwise distances, and if so, how? Before we discuss and clarify these issues, we discuss two qualitative justifications that were previously reported in the literature: the first based on a physical analogy, and the second using a statistical argument.

PMFs from the reversible work theorem

The first, qualitative justification of PMFs is due to Sippl, and based on an analogy with the statistical physics of liquids [37]. For liquids [8, 9, 13, 14, 37], the potential of mean force is related to the *pair correlation function* $g(r)$, which is given by:

$$g(r) = \frac{P(r)}{Q_R(r)}$$

where $P(r)$ and $Q_R(r)$ are the respective probabilities of finding two particles at a distance r from each other in the liquid and in the reference state. For liquids, the reference state is clearly defined; it corresponds to the ideal gas, consisting of non-interacting particles. The two-particle potential of mean force $W(r)$ is related to $g(r)$ by:

$$W(r) = -kT \log g(r) = -kT \log \frac{P(r)}{Q_R(r)} \quad (4)$$

According to the *reversible work theorem*, the two-particle potential of mean force $W(r)$ is the reversible work required to bring two particles in the liquid from infinite separation to a distance r from each other [8, 9].

Sippl justified the use of PMFs – a few years after he introduced them for use in protein structure prediction [7] – by appealing to the analogy with the reversible work theorem for liquids [37]. For liquids, $g(r)$ can be experimentally measured using small angle X-ray scattering; for proteins, $P(r)$ is obtained from the set of known protein structures, as explained in the previous section. The analogy described

above might provide some physical insight, but, as Ben-Naim writes in a seminal publication [13]: “the quantities, referred to as ‘statistical potentials,’ ‘structure based potentials,’ or ‘pair potentials of mean force,’ as derived from the protein data bank, are neither ‘potentials’ nor ‘potentials of mean force,’ in the ordinary sense as used in the literature on liquids and solutions.”

Another issue is that the analogy does not specify a suitable reference state for proteins. This is also reflected in the literature on statistical potentials; the construction of a suitable reference state continues to be an active research topic [3, 22, 38–41]. In the next section, we discuss a second, more recent justification that is based on probabilistic reasoning.

PMFs from likelihoods

Baker and co-workers [18] justified PMFs from a Bayesian point of view and used these insights in the construction of the coarse grained ROSETTA energy function; Samudrala and Moulton used similar reasoning for the RAPDF potential [42]. According to Bayesian probability calculus, the conditional probability $P(X | A)$ of a structure X , given the amino acid sequence A , can be written as:

$$P(X | A) = \frac{P(A | X) P(X)}{P(A)} \propto P(A | X) P(X)$$

$P(X | A)$ is proportional to the product of the likelihood $P(A | X)$ times the prior $P(X)$. By assuming that the likelihood can be approximated as a product of pairwise probabilities, and applying Bayes’ theorem, the likelihood can be written as:

$$P(A | X) \approx \prod_{i < j} P(a_i, a_j | r_{ij}) \propto \prod_{i < j} \frac{P(r_{ij} | a_i, a_j)}{P(r_{ij})} \quad (5)$$

where the product runs over all amino acid pairs a_i, a_j (with $i < j$), and r_{ij} is the distance between amino acids i and j . Obviously, the negative of the logarithm of expression (5) has the same functional form as the classic pairwise distance PMFs, with the denominator playing the role of the reference state in Eq. 1. The merit of this explanation is the qualitative demonstration that the functional form of a PMF can be obtained from probabilistic reasoning. Although this view is insightful – it rightfully drew the attention to the application of Bayesian methods to protein structure prediction – there is a more quantitative explanation, which does not rely on the incorrect assumption of pairwise decomposability [12–14, 43], and

leads to a different, *quantitative* conclusion regarding the nature of the reference state. This explanation is given in the next section.

A general statistical justification for PMFs

Expressions that resemble PMFs naturally result from the application of probability theory to solve a fundamental problem that arises in protein structure prediction: how to improve an imperfect probability distribution $Q(X)$ over a first variable X using a probability distribution $P(Y)$ over a second variable Y (see Fig. 2, Fig. 1 and Materials and Methods). We assume that Y is a deterministic function of X ; we write $Y(X)$ when necessary. In that case, X and Y are called *fine* and *coarse grained variables*, respectively. When Y is a function of X , the probability distribution $Q(X)$ automatically implies a probability distribution $Q(X, Y(X))$. This distribution has some unusual properties: $Q(X, Y(X)) = Q(X)$; and if $Y' \neq Y(X)$, it follows that $Q(X, Y') = 0$.

Typically, X represents *local* features of protein structure (such as backbone dihedral angles), while Y represents *nonlocal* features (such as hydrogen bonding, compactness or pairwise distances). However, the same reasoning also applies to other cases; for example, $P(Y)$ could represent information coming from experimental data, and $Q(X)$ could be embodied in an empirical force field as used in molecular mechanics [2, 44] (see Fig. 2).

Typically, the distribution $Q(X)$ in itself is not sufficient for protein structure prediction: it does not consider important nonlocal features such as hydrogen bonding, compactness or favorable amino acid interactions. As a result, $Q(X)$ is incorrect with respect to Y , and needs to be supplemented with a probability distribution $P(Y)$ that provides additional information. By construction, $P(Y)$ is assumed to be correct (or at least useful).

The above situation arises naturally in protein structure prediction. For example, $P(Y)$ could be a probability distribution over the radius of gyration, hydrogen bond geometry or the set of pairwise distances, and $Q(X)$ could be a fragment library [18] or a probabilistic model of local structure [45]. In Fig. 1, we used the example of a distribution over the radius of gyration for $P(Y)$ and a fragment library for $Q(X)$. Obviously, sampling from a fragment library and retaining structures with the desired nonlocal structure (radius of gyration, hydrogen bonding, etc.) is in principle possible, but in practice extremely inefficient.

How can $Q(X)$ be combined with $P(Y)$ in a meaningful way? As mentioned previously, simply

multiplying the two distributions – resulting in $P(Y(X))Q(X)$ – does not lead to the desired result as the two variables are obviously not independent. The correct solution follows from simple statistical considerations (see Materials and Methods), and is given by the following expression:

$$P(X) = \frac{P(Y(X))}{Q_R(Y(X))}Q(X) \quad (6)$$

We use the notation $P(X)$, as this distribution implies the desired distribution $P(Y)$ for $Y(X)$. The distribution $Q_R(Y)$ in the denominator is the probability distribution that is implied by $Q(X)$ over the coarse grained variable Y . Conceptually, dividing by $Q_R(Y)$ takes care of the signal in $Q(X)$ with respect to the coarse grained variable Y . The ratio in this expression corresponds to the probabilistic formulation of a PMF, and $Q_R(Y)$ corresponds to the reference state (see Materials and Methods).

In practice, $Q(X)$ is typically not evaluated directly, but brought in through conformational Monte Carlo sampling (see Materials and Methods); often sampling is based on a fragment library [18, 46], although other methods are possible, including sampling from a probabilistic model [45, 47, 48] or a suitable energy function [2, 44]. The ratio $P(Y)/Q_R(Y)$, which corresponds to the probabilistic formulation of a PMF, also naturally arises in the Markov chain Monte Carlo (MCMC) procedure (see Materials and Methods). An important insight is that, in this case, the conformational sampling method uniquely defines the reference state. Thus, in the case of a fragment library, the reference distribution $Q_R(Y)$ is the probability distribution over Y that is obtained by sampling conformations solely using the fragment library.

As the method we have introduced here invariably relies on the ratio of two probability distributions – one regarding protein structure and the other regarding a well-defined reference state – we refer to it as the *reference ratio method*. In the next section, we show that the standard pairwise distance PMFs can be seen as an approximation of the reference ratio method.

Pairwise distance PMFs explained

In this section, we apply the reference ratio method to the standard, pairwise distance case. In the classic PMF approach, one considers the vector of pairwise distances R between the amino acids. In this case,

it is usually assumed that we can write

$$P(R | A) \propto \prod_{i < j} P(r_{ij} | a_i, a_j) \quad (7)$$

where the product runs over all amino acid pairs a_i, a_j (with $i < j$), and r_{ij} is their matching distance. Clearly, the assumption that the joint probability can be written as a product of pairwise probabilities is not justified [12, 13, 43], but in practice this assumption often provides useful results [22]. In order to obtain protein-like conformations, $P(R | A)$ needs to be combined with an appropriate probability distribution $Q(X | A)$ that addresses the local features of the polypeptide chain. Applying Eq. 6 to this case results in the following expression:

$$P(X | A) \propto \frac{\prod_{i < j} P(r_{ij} | a_i, a_j)}{\prod_{i < j} Q_R(r_{ij} | a_i, a_j)} Q(X | A)$$

where the denominator $Q_R(\cdot)$ is the probability distribution over the pairwise distances as induced by the distribution $Q(X | A)$. The ratio in this expression corresponds to the probabilistic expression of a PMF. The reference state is thus determined by $Q(X | A)$: it reflects the probability of generating a set of pairwise distances using local structure information alone. Obviously, as $Q(X | A)$ is conditional upon the amino acid sequence A , the reference state becomes sequence dependent as well.

We again emphasize that the assumption of pairwise decomposability in Eq. 7 is incorrect [12–14, 43]. Therefore, the application of the reference ratio method results in a useful approximation, at best. As a result, the optimal definition of the reference state also needs to compensate for the errors implied by the invalid assumption. As is well established that distance dependent PMFs perform well with a suitable definition of the reference state [3, 22, 38–40], and the incorrect pairwise decomposability assumption impairs a rigorous statistical analysis, we do not discuss this type of PMFs further. Indeed, for pairwise distance PMFs, the main challenge lies in developing better probabilistic models of sets of pairwise distances [49].

The pairwise distance PMFs currently used in protein structure prediction are thus not statistically rigorous, because they do not make use of a proper joint probability distribution over the pairwise distances, which are strongly intercorrelated due to the connectivity of molecules. A rigorous application of the reference ratio method would require the construction of a proper joint probability distribution over

pairwise distances. This is certainly possible in principle, but currently, as far as we know, a challenging open problem and beyond the scope of this article. However, we have clarified that the idea of using a reference state is correct and valid, and that this state has a very precise definition. Therefore, in the next two sections, we show instead how statistically valid quantities, similar to PMFs, can be obtained for very different coarse grained variables.

A generalized PMF: radius of gyration

As a first application of the reference ratio method, we consider the task of sampling protein conformations with a given probability distribution $P(r_g)$ for the radius of gyration r_g . For $P(r_g)$, we chose a Gaussian distribution with mean $\mu = 22 \text{ \AA}$ and standard deviation $\sigma = 2 \text{ \AA}$. This choice is completely arbitrary; it simply serves to illustrate that the reference ratio method allows imposing an exact probability distribution over a certain feature of interest. Applying Eq. 6 results in:

$$P(X | A) = \frac{P(r_g(X))}{Q_R(r_g(X) | A)} Q(X | A) \quad (8)$$

For $Q(X | A)$, we used TorusDBN – a graphical model that allows sampling of plausible backbone angles [45] – and sampled conditional on the amino acid sequence A of ubiquitin (see Materials and Methods). $Q_R(r_g | A)$ is the probability distribution of the radius of gyration for structures sampled solely from TorusDBN, which was determined using generalized multihistogram MCMC sampling (see Materials and Methods).

In Fig. 3, we contrast sampling from Eq. 8 with sampling from $P(r_g(X))Q(X | A)$. In the latter case, the reference state is not properly taken into account, which results in a significant shift towards higher radii of gyration. In contrast, the distribution of r_g for the correct distribution $P(X)$, given by Eq. 8, is indistinguishable from the target distribution. This qualitative result is confirmed by the Kullback-Leibler divergence [50] – a natural distance measure for probability distributions expressed in bits – between the target distribution and the resulting marginal distributions of r_g . Adding $Q_R(r_g(X) | A)$ to the denominator diminishes the distance from 0.08 to 0.001 bits. For this particular PMF, the effect of using the correct reference state is significant, but relatively modest; in the next section, we discuss an application where its effect is much more pronounced.

Iterative optimization of PMFs: hydrogen bonding

Here, we demonstrate that PMFs can be optimized iteratively, which is particularly useful if the reference probability distribution $Q_R(Y | A)$ is difficult to estimate. We illustrate the method with a target distribution that models the hydrogen bonding network using a multinomial distribution.

We describe the hydrogen bonding network (H) with eight integers (for details, see Materials and Methods). Three integers (n_α, n_β, n_c) represent the number of residues that do not partake in hydrogen bonds in α -helices, β -sheets and coils, respectively. The five remaining integers ($n_{\alpha\alpha}, n_{\beta\beta}, n_{cc}, n_{\alpha c}, n_{\beta c}$) represent the number of hydrogen bonds within α -helices, within β -strands, within coils, between α -helices and coils, and between β -strands and coils, respectively.

As target distribution $P(H)$ over these eight integers, we chose a multinomial distribution whose parameters were derived from the native structure of protein G (see Materials and Methods). $P(H)$ provides information, regarding protein G, on the number of hydrogen bonds and the secondary structure elements involved, but does not specify *where* the hydrogen bonds or secondary elements occur. As in the previous section, we use TorusDBN as the sampling distribution $Q(X | A)$; we sample backbone angles conditional on the amino acid sequence A of protein G. Native secondary structure information was *not* used in sampling from TorusDBN.

The reference distribution $Q_R(H | A)$, due to TorusDBN, is very difficult to estimate correctly for several reasons: its shape is unknown and presumably complex; its dimensionality is high; and the data is very sparse with respect to β -sheet content. Therefore, $Q_R(H | A)$ can only be approximated, which results in a suboptimal PMF. A key insight is that one can apply the method iteratively until a satisfactory PMF is obtained (see Fig. 2, dashed line). In each iteration, the (complex) reference distribution is approximated using a simple probability distribution; we illustrate the method by using a multinomial distribution, whose parameters are estimated by maximum likelihood estimation in each iteration, using the conformations generated in the previous iteration. In the first iteration, we simply set the reference distribution equal to the uniform distribution.

Formally, the procedure works as follows. In iteration $i + 1$, the distribution $P_i(H | A)$ is improved using the samples generated in iteration i :

$$P_{i+1}(X | A) = \frac{P(H(X))}{P_{R,i}(H(X) | A)} P_i(X | A) \quad (9)$$

where $P_{R,i}(H | A)$ is the reference distribution estimated from the samples generated in the i -th iteration, $P_0(X) = Q(X | A)$ stems from TorusDBN, and $P_{R,0}(H | A)$ is the uniform distribution. After each iteration, the set of samples is enriched in hydrogen bonds, and the reference distribution $P_{R,i}(H | A)$ can be progressively estimated more precisely. Note that in the first iteration, we simply use the product of the target and the sampling distribution; no reference state is involved.

Fig. 4 shows the evolution of the fractions versus the iteration number for the eight hydrogen bond categories; the structures with minimum energy for all six iterations are shown in Fig. 5. In the first iteration, the structure with minimum energy (highest probability) consists of a single α -helix; β -sheets are entirely absent (see Fig. 5, structure 1). Already in the second iteration, β -strands start to pair, and in the third and higher iterations complete sheets are readily formed. The iterative optimization of the PMF quickly leads to a dramatic enrichment in β -sheet structures, as desired, and the fractions of the eight categories become very close to the native values (Fig. 4).

Conclusions

The strengths and weaknesses of PMFs can be rigorously explained based on simple probabilistic considerations, which leads to some surprising new insights of direct practical relevance. First, we have made clear that PMFs naturally arise when two probability distributions need to be combined in a meaningful way. One of these distributions typically addresses local structure, and its contribution often arises from conformational sampling. Each conformational sampling method thus requires its own reference state and corresponding reference distribution; this is likely the main reason behind the large number of different reference states reported in the literature [3, 22, 38–41]. If the sampling method is conditional upon the amino acid sequence, the reference state necessarily also depends on the amino acid sequence.

Second, conventional applications of pairwise distance PMFs usually lack two necessary features to make them fully rigorous: the use of a proper probability distribution over pairwise distances in proteins for $P(Y | A)$, and the recognition that the reference state is rigorously defined by the conformational sampling scheme used, that is, $Q(X | A)$. Usually, the reference state is derived from external physical considerations [11, 51].

Third, PMFs are not tied to pairwise distances, but generalize to any coarse grained variable. Attempts to develop similar quantities that, for example, consider solvent exposure [52, 53], relative side chain orientations [54], backbone dihedral angles [55, 56] or hydrogen bonds [37] are thus, in principle, en-

tirely justified. Hence, our probabilistic interpretation opens up a wide range of possibilities for advanced, well-justified energy functions based on sound probabilistic reasoning; the main challenge is to develop proper probabilistic models of the features of interest and the estimation of their parameters [49, 57]. Strikingly, the example applications involving radius of gyration and hydrogen bonding that we presented in this article *are* statistically valid and rigorous, in contrast to the traditional pairwise distance PMFs.

Finally, our results reveal a straightforward way to optimize PMFs. Often, it is difficult to estimate the probability distribution that describes the reference state. In that case, one can start with an approximate PMF, and apply the method iteratively. In each iteration, a new reference state is estimated, with a matching probability distribution. In that way, one iteratively attempts to sculpt an energy funnel [24, 32–36]. We illustrated this approach with a probabilistic model of the hydrogen bond network. Although iterative application of the inverse Boltzmann formula has been described before [24, 35, 58, 59], its theoretical justification, optimal definition of the reference state and scope remained unclear.

As the traditional pairwise distance PMFs used in protein structure prediction arise from the imperfect application of a statistically valid and rigorous procedure with a much wider scope, we consider it highly desirable that the name “potential of mean force” should be reserved for true, physically valid quantities [13]. Because the statistical quantities we discussed invariably rely on the use of a ratio of two probability distributions, one concerning protein structure and the other concerning the (now well defined) reference state, we suggest the name “reference ratio distribution” deriving from the application of the “reference ratio method”.

Pairwise distance PMFs, as used in protein structure prediction, are not physically justified potentials of mean force or free energies [2, 13] and the reference state does not depend on external physical considerations; the same is of course true for our generalization. However, these PMFs are approximations of statistically valid and rigorous quantities, and these quantities can be generalized beyond pairwise distances to other aspects of protein structure. The fact that these quantities are not potentials of mean force or free energies is of no consequence for their statistical rigor or practical importance – both of which are considerable. Our results thus vindicate, formalize and generalize Sippl’s original and seminal idea [7]. After about twenty years of controversy, PMFs – or rather the statistical quantities that we have introduced in this article – are ready for new challenges.

Materials and methods

Outline of the problem

We consider a joint probability distribution $Q(X, Y)$ and a probability distribution $P(Y)$ over two variables of interest, X and Y , where Y is a deterministic function of X ; we write $Y(X)$ when relevant. Note that because Y is a function of X , it follows that $Q(X) = Q(X, Y(X))$; and if $Y' \neq Y(X)$, then $Q(X, Y') = 0$.

We assume that $P(Y)$ is a meaningful and informative distribution for Y . Next, we note that $Q(X, Y)$ implies a matching marginal probability distribution $Q_R(Y)$ (where the subscript R refers to the fact that $Q_R(Y)$ corresponds to the reference state, as we will show below):

$$Q_R(Y) = \int Q(X, Y) dX$$

We consider the case where $Q_R(Y)$ differs substantially from $P(Y)$; hence, $Q_R(Y)$ can be considered as incorrect. On the other hand, we also assume that the conditional distribution $Q(X | Y)$ is indeed meaningful and informative (see next section). This distribution is given by:

$$Q(X | Y) = \begin{cases} 0 & \text{if } Y \neq Y(X) \\ \frac{Q(X)}{\int Q(X') \delta(Y(X') - Y) dX'} & \text{if } Y = Y(X) \end{cases} \quad (10)$$

where $\delta(\cdot)$ is the delta function. The question is now how to combine the two distributions $P(Y)$ and $Q(X)$ – each of which provide useful information on X and Y – in a meaningful way. Before we provide the solution, we illustrate how this problem naturally arises in protein structure prediction.

Application to protein structure

In protein structure prediction, $Q(X, Y)$ is often embodied in a fragment library; in that case, X is a set of atomic coordinates obtained from assembling a set of polypeptide fragments. Of course, $Q(X, Y)$ could also arise from a probabilistic model, a pool of known protein structures, or any other conformational sampling method. The variable Y could, for example, be the radius of gyration, the hydrogen bond network or the set of pairwise distances. If Y is a deterministic function of X , the two variables are called

coarse grained and *fine grained* variables, respectively. For example, sampling a set of dihedral angles for the protein backbone uniquely defines the hydrogen bond geometry between any of the backbone atoms.

Above, we assumed that $Q(X | Y)$ is a meaningful distribution. This is often a reasonable assumption; fragment libraries, for example, originate from real protein structures, and conditioning on protein-like compactness or hydrogen bonding will thus result in a meaningful distribution. Of course, sampling solely from $Q(X, Y)$ is not an efficient strategy to obtain hydrogen bonded or compact conformations, as they will be exceedingly rare. We now provide the solution of the problem outlined in the previous section, and discuss its relevance to the construction of PMFs.

Solution for a proper joint distribution

A first step on the way to the solution is to note that the product rule of probability theory allows us to write:

$$P(X, Y) = P(Y)P(X | Y)$$

As only $P(Y)$ is given, we need to make a reasonable choice for $P(X | Y)$. We assume, as discussed before, that $Q(X | Y)$ is a meaningful choice, which leads to:

$$P(X, Y) = P(Y)Q(X | Y)$$

In the next step, we apply the product formula of probability theory to the second factor $Q(X | Y)$, and obtain:

$$P(X, Y) = P(Y) \frac{Q(X, Y)}{Q_R(Y)} \tag{11}$$

The distribution $P(X, Y)$ has the correct marginal distribution $P(Y)$.

In the next two sections, we discuss how this straightforward result can be used to great advantage for understanding and generalizing PMFs. First, we show that the joint distribution specified by Eq. 11 can be reduced to a surprisingly simple functional form. Second, we discuss how this result can be used in MCMC sampling. In both cases, expressions that correspond to a PMF arise naturally.

PMFs from combining distributions

Using the product rule of probability theory, Eq. 11 can be written as:

$$P(X, Y) = P(Y) \frac{Q(Y | X) Q(X)}{Q_R(Y)}$$

Because the coarse grained variable Y is a deterministic function of the fine grained variable X , $Q(Y | X)$ is the delta function:

$$P(X, Y) = P(Y) \frac{\delta(Y - Y(X)) Q(X)}{Q_R(Y)} \quad (12)$$

Finally, we integrate out the, now redundant, coarse grained variable Y from the expression:

$$\begin{aligned} P(X) &= \int P(X, Y) dY \\ &= \int P(Y) \frac{\delta(Y - Y(X)) Q(X)}{Q_R(Y)} dY \\ &= \frac{P(Y(X))}{Q_R(Y(X))} Q(X) \end{aligned}$$

and obtain our central result (Eq. 6). Sampling from $P(X)$ will result in the desired marginal probability distribution $P(Y)$. The influence of the fine grained distribution $Q(X, Y)$ is apparent in the fact that $P(X | Y)$ is equal to $Q(X | Y)$. The ratio in this expression corresponds to the usual probabilistic formulation of a PMF; the distribution $Q_R(Y)$ corresponds to the reference state. In the next section, we show that PMFs also naturally arise when $P(Y)$ and $Q(X, Y)$ are used together in Metropolis-Hastings sampling.

PMFs from Metropolis-Hastings sampling

Here, we show that Metropolis-Hastings sampling from the distribution specified by Eq. 11, using $Q(X, Y)$ as a proposal distribution, naturally results in expressions that are equivalent to PMFs. The derivation is also valid if the proposal distribution depends on the previous state, provided $Q(X, Y)$ satisfies the detailed balance condition.

According to the standard Metropolis-Hastings method [60], one can sample from a probability distribution $\Pi(X, Y)$ by generating a Markov chain where each state X', Y' depends only on the previous

state X, Y . The new state X', Y' is generated using a proposal distribution $\pi(Y', X' | Y, X)$, which includes $\pi(X', Y' | X, Y) = \pi(X', Y')$ as a special case. According to the Metropolis-Hastings method, the proposal X', Y' is accepted with a probability α :

$$\alpha(X', Y' | X, Y) = \min(1, p),$$

$$p = \frac{\Pi(X', Y')}{\Pi(X, Y)} \times \frac{\pi(X, Y | X', Y')}{\pi(X', Y' | X, Y)} \quad (13)$$

where Y, X is the starting state, and Y', X' is the next proposed state. We assume that the proposal distribution $\pi(X', Y' | X, Y)$ satisfies the detailed balance condition:

$$\pi(X', Y' | X, Y)\pi(X, Y) = \pi(X, Y | X', Y')\pi(X', Y')$$

As a result, we can always write Eq. 13 as:

$$\frac{\Pi(X', Y')}{\Pi(X, Y)} \times \frac{\pi(X, Y)}{\pi(X', Y')}$$

The Metropolis-Hastings expression (Eq. 13), applied to the distribution specified by Eq. 11 and using $Q(X', Y')$ or $Q(X', Y' | X, Y)$ as the proposal distribution, results in:

$$\frac{P(Y')Q_R(Y)Q(X', Y')}{P(Y)Q_R(Y')Q(X, Y)} \times \frac{Q(X, Y)}{Q(X', Y')}$$

which reduces to:

$$\frac{P(Y')}{P(Y)} \times \frac{Q_R(Y)}{Q_R(Y')} \quad (14)$$

Hence, we see that the Metropolis-Hastings method requires the evaluation of ratios of the form $P(Y)/Q_R(Y)$ when $Q(X', Y')$ or $Q(X', Y' | X, Y)$ is used as the proposal distribution; these ratios correspond to the usual probabilistic formulation of a PMF. Finally, when Y is a deterministic function of X , the proposal distribution reduces to $Q(X')$ or $Q(X' | X)$, and Eq. 14 becomes:

$$\frac{P(Y(X'))}{P(Y(X))} \times \frac{Q_R(Y(X))}{Q_R(Y(X'))}$$

Application to radius of gyration and hydrogen bonding

Conformational sampling from a suitable $Q(X | A)$ was done using TorusDBN [45] as implemented in Phaistos [61]; backbone angles (ϕ, ψ and ω) were sampled conditional on the amino acid sequence. We used standard fixed bond lengths and bond angles in constructing the backbone coordinates from the angles, and represented all side chains (except glycine and alanine) with one dummy atom with a fixed position [61].

For the radius of gyration application, we first determined $Q_R(r_g | A)$ using the multi-canonical MCMC method to find the sampling weights $w(r_g)$ that yield a flat histogram [62]. Sampling from the resulting joint distribution (Eq. 8) was done using the same method. In both cases, we used 50 million iterations; the r_g bin size was 0.08 Å. Sampling from TorusDBN was done conditional on the amino acid sequence A of ubiquitin (76 residues, PDB code 1UBQ).

For the hydrogen bond application, sampling from the PMFs was done in the $1/k$ -ensemble [63], using the Metropolis-Hastings algorithm and the generalized multihistogram method for updating the weights [62]. In each iteration i , 50,000 samples (out of 50 million Metropolis-Hastings steps) were generated, and the parameters of the multinomial distribution $Q_{R,i}(H)$ were subsequently obtained using maximum likelihood estimation. Hydrogen bonds were defined as follows: the N, O distance is below 3.5 Å, and the angles formed by O, H, N and C, O, H are both greater than 100° . Each carbonyl group was assumed to be involved in at most one hydrogen bond; in case of multiple hydrogen bond partners, the one with the lowest H, O distance was selected. Each residue was assigned to one of the eight possible hydrogen bond categories ($n_\alpha, n_\beta, n_c, n_{\alpha\alpha}, n_{\beta\beta}, n_{cc}, n_{\alpha c}, n_{\beta c}$) based on the presence of hydrogen bonding at its carbonyl group and the secondary structure assignments (for both bond partners) by TorusDBN. The target distribution – the multinomial distribution $P(H)$ used in Eq. 9 – was obtained by maximum likelihood estimation using the number of hydrogen bonds, for all eight categories, in the native structure of protein G (56 residues, PDB code 2GB1). Sampling from TorusDBN was done conditional on the amino acid sequence of protein G; native secondary structure information was *not* used.

Author contributions

T.H., M.B. and M.P. are joint first authors. T.H. and J.F.B. developed the theory. M.B. and M.P. performed the simulations. The remaining authors contributed new tools. T.H. wrote the paper.

Acknowledgments

We acknowledge funding by the *Danish Program Commission on Nanoscience, Biotechnology and IT* (NaBiIT, project: “Simulating proteins on a millisecond time-scale”, 2106-06-0009), the *Danish Research Council for Technology and Production Sciences* (FTP, project: “Protein structure ensembles from mathematical models”, 274-09-0184) and the *Danish Council for Independent Research* (FNU, project: “A Bayesian approach to protein structure determination”, 272-08-0315).

References

1. Anfinsen CB (1973) Principles that govern the folding of protein chains. *Science* 181: 223–230.
2. Moult J (1997) Comparison of database potentials and molecular mechanics force fields. *Curr Opin Struct Biol* 7: 194–199.
3. Shen MY, Sali A (2006) Statistical potential for assessment and prediction of protein structures. *Protein Sci* 15: 2507–2524.
4. Tanaka S, Scheraga HA (1976) Medium- and long-range interaction parameters between amino acids for predicting three-dimensional structures of proteins. *Macromolecules* 9: 945–950.
5. Miyazawa S, Jernigan R (1985) Estimation of effective interresidue contact energies from protein crystal structures: quasi-chemical approximation. *Macromolecules* 18: 534–552.
6. Miyazawa S, Jernigan R (1999) An empirical energy potential with a reference state for protein fold and sequence recognition. *Proteins* 36: 357–369.
7. Sippl MJ (1990) Calculation of conformational ensembles from potentials of mean force. An approach to the knowledge-based prediction of local structures in globular proteins. *J Mol Biol* 213: 859–883.
8. Chandler D (1987) *Introduction to Modern Statistical Mechanics*. Oxford University Press, USA.
9. McQuarrie D (2000) *Statistical mechanics*. University Science Books, USA.
10. Finkelstein A, Badretdinov A, Gutin A (1995) Why do protein architectures have Boltzmann-like statistics? *Proteins Struct Func Gen* 23: 142–150.

11. Rooman M, Wodak S (1995) Are database-derived potentials valid for scoring both forward and inverted protein folding? *Protein Eng* 8: 849-858.
12. Thomas PD, Dill KA (1996) Statistical potentials extracted from protein structures: how accurate are they? *J Mol Biol* 257: 457-469.
13. Ben-Naim A (1997) Statistical potentials extracted from protein structures: Are these meaningful potentials? *J Chem Phys* 107: 3698-3706.
14. Koppensteiner WA, Sippl MJ (1998) Knowledge-based potentials—back to the roots. *Biochemistry Mosc* 63: 247-252.
15. Shortle D (2003) Propensities, probabilities, and the Boltzmann hypothesis. *Protein Sci* 12: 1298-1302.
16. Kirtay C, Mitchell J, Lumley J (2005) Knowledge based potentials: The reverse Boltzmann methodology, virtual screening and molecular weight dependence. *QSAR & Combinatorial Sci* 24: 527-536.
17. Muegge I (2006) PMF scoring revisited. *J Med Chem* 49: 5895-5902.
18. Simons KT, Kooperberg C, Huang E, Baker D (1997) Assembly of protein tertiary structures from fragments with similar local sequences using simulated annealing and Bayesian scoring functions. *J Mol Biol* 268: 209-225.
19. Colubri A, Jha A, Shen M, Sali A, Berry R, et al. (2006) Minimalist representations and the importance of nearest neighbor effects in protein folding simulations. *J Mol Biol* 363: 835-857.
20. Sippl MJ (1993) Recognition of errors in three-dimensional structures of proteins. *Proteins* 17: 355-362.
21. Eramian D, Shen M, Devos D, Melo F, Sali A, et al. (2006) A composite score for predicting errors in protein structure models. *Protein Sci* 15: 1653-1666.
22. Rykunov D, Fiser A (2010) New statistical potential for quality assessment of protein models and a survey of energy functions. *BMC Bioinformatics* 11: 128.
23. Jones DT, Taylor WR, Thornton JM (1992) A new approach to protein fold recognition. *Nature* 358: 86-89.

24. Májek P, Elber R (2009) A coarse-grained potential for fold recognition and molecular dynamics simulations of proteins. *Proteins* 76: 822–836.
25. Gohlke H, Hendlich M, Klebe G (2000) Knowledge-based scoring function to predict protein-ligand interactions1. *J Mol Biol* 295: 337–356.
26. Gilis D, Rooman M (1997) Predicting protein stability changes upon mutation using database-derived potentials: solvent accessibility determines the importance of local versus non-local interactions along the sequence1. *J Mol Biol* 272: 276–290.
27. Gilis D, Rooman M (2000) PoPMuSiC, an algorithm for predicting protein mutant stability changes. Application to prion proteins. *Protein Eng* 13: 849–856.
28. Su Y, Zhou A, Xia X, Li W, Sun Z (2009) Quantitative prediction of protein-protein binding affinity with a potential of mean force considering volume correction. *Protein Sci* 18: 2550–2558.
29. Chandler D (2005) Interfaces and the driving force of hydrophobic assembly. *Nature* 437: 640–647.
30. Bowman GR, Pande VS (2009) Simulated tempering yields insight into the low-resolution Rosetta scoring functions. *Proteins* 74: 777–788.
31. Shmygelska A, Levitt M (2009) Generalized ensemble methods for de novo structure prediction. *Proc Natl Acad Sci U S A* 106: 1415–1420.
32. Bryngelson J, Wolynes P (1987) Spin glasses and the statistical mechanics of protein folding. *Proc Natl Acad Sci U S A* 84: 7524–7528.
33. Leopold P, Montal M, Onuchic J (1992) Protein folding funnels: a kinetic approach to the sequence-structure relationship. *Proc Natl Acad Sci U S A* 89: 8721–8725.
34. Dill K, Chan H (1997) From Levinthal to pathways to funnels. *Nat Struct Biol* 4: 10–19.
35. Reith D, Pütz M, Müller-Plathe F (2003) Deriving effective mesoscale potentials from atomistic simulations. *J Comput Chem* 24: 1624–1636.
36. Fain B, Levitt M (2003) Funnel sculpting for in silico assembly of secondary structure elements of proteins. *Proc Natl Acad Sci U S A* 100: 10700–10705.

37. Sippl MJ, Ortner M, Jaritz M, Lackner P, Flockner H (1996) Helmholtz free energies of atom pair interactions in proteins. *Fold Des* 1: 289–98.
38. Zhang C, Liu S, Zhou H, Zhou Y (2004) An accurate, residue-level, pair potential of mean force for folding and binding based on the distance-scaled, ideal-gas reference state. *Protein Sci* 13: 400–411.
39. Cheng J, Pei J, Lai L (2007) A free-rotating and self-avoiding chain model for deriving statistical potentials based on protein structures. *Biophys J* 92: 3868–3877.
40. Rykunov D, Fiser A (2007) Effects of amino acid composition, finite size of proteins, and sparse statistics on distance-dependent statistical pair potentials. *Proteins* 67: 559–568.
41. Bernard B, Samudrala R (2008) A generalized knowledge-based discriminatory function for biomolecular interactions. *Proteins* 76: 115–128.
42. Samudrala R, Moult J (1998) An all-atom distance-dependent conditional probability discriminatory function for protein structure prediction. *J Mol Biol* 275: 895–916.
43. Pearl J (1988) Probabilistic reasoning in intelligent systems, Morgan Kaufmann, San Francisco, USA, chapter 3. pp. 108–115.
44. Lazaridis T, Karplus M (2000) Effective energy functions for protein structure prediction. *Curr Opin Struct Biol* 10: 139–145.
45. Boomsma W, Mardia KV, Taylor CC, Ferkinghoff-Borg J, Krogh A, et al. (2008) A generative, probabilistic model of local protein structure. *Proc Natl Acad Sci U S A* 105: 8932–8937.
46. Sippl M, Hendlich M, Lackner P (1992) Assembly of polypeptide and protein backbone conformations from low energy ensembles of short fragments: Development of strategies and construction of models for myoglobin, lysozyme, and thymosin β 4. *Protein Sci* 1: 625–640.
47. Hamelryck T, Kent J, Krogh A (2006) Sampling realistic protein conformations using local structural bias. *PLoS Comput Biol* 2: e131.
48. Zhao F, Peng J, DeBartolo J, Freed K, Sosnick T, et al. (2010) A probabilistic and continuous model of protein conformational space for template-free modeling. *J Comput Biol* 17: 783–798.

49. Hamelryck T (2009) Probabilistic models and machine learning in structural bioinformatics. *Stat Methods Med Res* 18: 505–526.
50. Kullback S, Leibler R (1951) On information and sufficiency. *Annals Math Stat* 22: 79–86.
51. Zhou H, Zhou Y (2002) Distance-scaled, finite ideal-gas reference state improves structure-derived potentials of mean force for structure selection and stability prediction. *Protein Sci* 11: 2714–2726.
52. Bowie J, Luthy R, Eisenberg D (1991) A method to identify protein sequences that fold into a known three-dimensional structure. *Science* 253: 164–164.
53. Liithy R, Bowie J, Eisenberg D (1992) Assessment of protein models with three-dimensional profiles. *Nature* 356: 83–85.
54. Buchete NV, Straub JE, Thirumalai D (2004) Development of novel statistical potentials for protein fold recognition. *Curr Opin Struct Biol* 14: 225–232.
55. Rooman M, Kocher J, Wodak S (1991) Prediction of protein backbone conformation based on seven structure assignments: Influence of local interactions. *J Mol Biol* 221: 961–979.
56. Kocher J, Rooman M, Wodak S (1994) Factors influencing the ability of knowledge-based potentials to identify native sequence-structure matches. *J Mol Biol* 235: 1598–1613.
57. Simons KT, Ruczinski I, Kooperberg C, Fox BA, Bystroff C, et al. (1999) Improved recognition of native-like protein structures using a combination of sequence-dependent and sequence-independent features of proteins. *Proteins* 34: 82–95.
58. Thomas P, Dill K (1996) An iterative method for extracting energy-like quantities from protein structures. *Proc Natl Acad Sci U S A* 93: 11628–11633.
59. Huang S, Zou X (2006) An iterative knowledge-based scoring function to predict protein-ligand interactions: I. Derivation of interaction potentials. *J Comp Chem* 27: 1866–1875.
60. Gilks W, Richardson S, Spiegelhalter D (1996) Markov chain Monte Carlo in practice. Chapman & Hall/CRC, USA.
61. Borg M, Mardia K, Boomsma W, Frellsen J, Harder T, et al. (2009) A probabilistic approach to protein structure prediction: PHAISTOS in CASP9. In: Gusnanto A, Mardia K, Fallaize C,

- editors, LASR2009 - Statistical tools for challenges in bioinformatics. Leeds University Press, Leeds, UK, pp. 65–70.
62. Ferkinghoff-Borg J (2002) Optimized Monte Carlo analysis for generalized ensembles. *Eur Phys J B* 29: 481-484.
63. Hesselbo B, Stinchcombe R (1995) Monte Carlo simulation and global optimization without parameters. *Phys Rev Lett* 74: 2151–2155.
64. Delano WL (2002) The PyMOL Molecular Graphics System. Palo Alto, CA, USA: DeLano Scientific.

Figure legends




Figure 1. Illustration of the central idea presented in this article. In this example, the goal is to sample conformations with a given distribution $P(r_g)$ for the radius of gyration r_g , and a plausible local structure. $P(r_g)$ could, for example, be derived from known structures in the Protein Data Bank (PDB, left box). $Q(X)$ is a probability distribution over local structure X , typically embodied in fragment library (right box). In order to combine $Q(X)$ and $P(r_g)$ in a meaningful way (see text), the two distributions are multiplied and divided by $Q_R(r_g)$ (formula at the bottom); $Q_R(r_g)$ is the probability distribution over the radius of gyration for conformations sampled solely from the fragment library (that is, $Q(X)$). The probability distribution $P(X)$ will generate conformations with plausible local structures (due to $Q(X)$), while their radii of gyration will be distributed according to $P(r_g)$, as desired. This simple idea lies at the theoretical heart of the PMF expressions used in protein structure prediction.

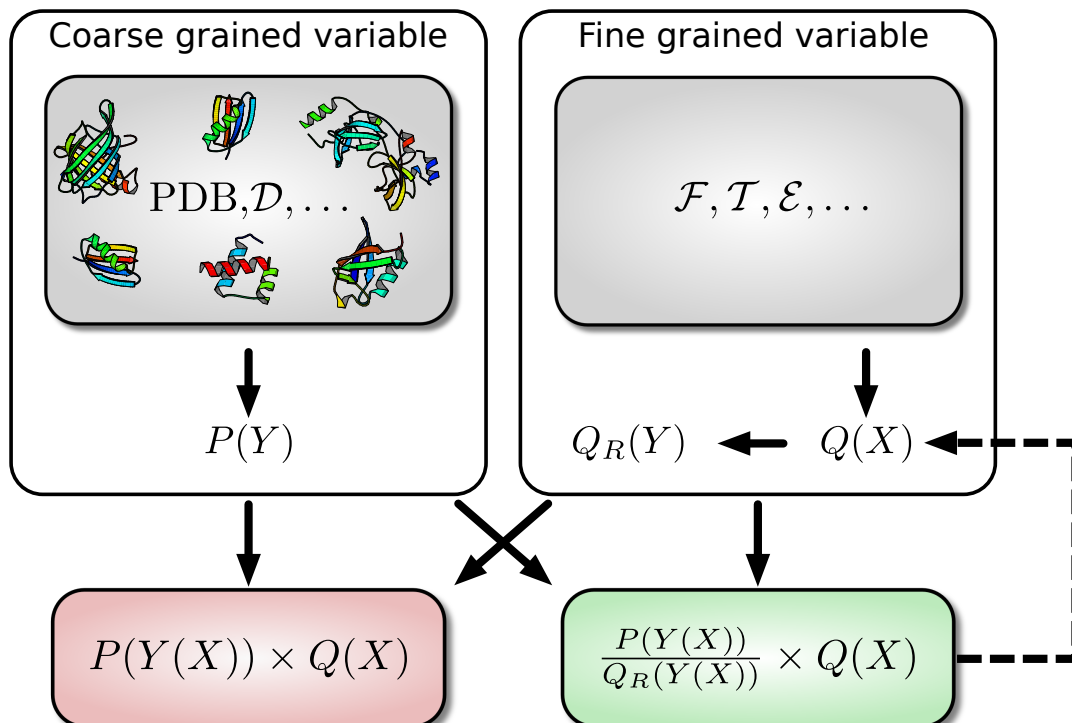


Figure 2. General statistical justification of PMFs. The goal is to combine a distribution $Q(X)$ over a fine grained variable X (top right), with a probability distribution $P(Y)$ over a coarse grained variable $Y(X)$ (top left). $Q(X)$ could be, for example, embodied in a fragment library (\mathcal{F}), a probabilistic model of local structure (\mathcal{T}) or an energy function (\mathcal{E}); Y could be, for example, the radius of gyration, the hydrogen bond network, or the set of pairwise distances. $P(Y)$ usually reflects the distribution of Y in known protein structures (PDB), but could also stem from experimental data (\mathcal{D}). Sampling from $Q(X)$ results in a distribution $Q_R(Y)$ that differs from $P(Y)$. Multiplying $P(Y)$ and $Q(X)$ does not result in the desired distribution for Y either (red box); the correct result requires dividing out the signal with respect to Y due to $Q(X)$ (green box). The *reference* distribution $Q_R(Y)$ in the denominator corresponds to the contribution of the reference state in a PMF. If $Q_R(Y)$ is only approximately known, the method can be applied iteratively (dashed arrow). In that case, one attempts to iteratively sculpt an energy funnel. The procedure is statistically rigorous provided $Q(X)$ and $P(Y)$ are proper probability distributions; this is usually not the case for conventional pairwise distance PMFs.




Figure 3. A PMF based on the radius of gyration. The goal is to adapt a distribution $Q(X | A)$ – which allows sampling of local structures – such that a given target distribution $P(r_g)$ is obtained. For A , we used the amino acid sequence of ubiquitin. Sampling from $Q(X | A)$ alone results in a distribution with an average r_g of about 27 Å (triangles). Sampling using the correct expression (open circles), given by Eq. 8, results in a distribution that coincides with the target distribution (solid line). Not taking the reference state into account results in a significant shift towards higher r_g (black circles).




Figure 4. Iterative estimation of a PMF. For each of the eight hydrogen bond categories (see text), the black bar to the right denotes the fraction of occurrence $f(n)$ in the native structure of protein G. The gray bars denote the fractions of the eight categories in samples from each iteration; the first iteration is shown to the left in light gray. In the last iteration (iteration 6; dark gray bars, right) the values are very close to the native values for all eight categories. Note that hydrogen bonds between β -strands are nearly absent in the first iteration (category $n_{\beta\beta}$).




Figure 5. Highest probability structures for each iteration. The structures with highest probability out of 50,000 samples for all six iterations (indicated by a number) are shown as cartoon representations. The N-terminus is shown in blue. The figure was made using PyMOL [64].